Impact of I/O and Data Management in Ensemble Large Scale Climate Forecasting Using EC-Earth3

Muhammad Asif[1], Andrés Cencerrado[2], Oriol Mula-Valls[1], Domingo Manubens[1], Francisco Doblas-Reyes[1, 3] and Ana Cortés[2]

(1): Institut Català de Ciències del Clima (IC3), Barcelona, Spain. (2): Computer Architecture and Operating Systems Department, Universitat Autònoma de Barcelona, 08193 Bellaterra, Barcelona, Spain. (3): Institució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain. Contact: `muhammad.asif@ic3.cat`

Abstract

The EC-Earth climate model is a seamless Earth System Model (ESM) used to carry out climate research in 24 academic institutions and meteorological services from 11 countries in Europe. This model couples several components and it is continuously under development.

In this work we present a study regarding the impact of the I/O and data management when using EC-Earth in well-known supercomputing environments.

Most large-scale and long-term climate simulators have been developed bearing in mind the paramount importance of its scalability. However, the computational capabilities of the *High Performance Computing* (HPC) environments increase at so great speed that it is almost impossible to re-implement the whole models so that they are able to exploit efficiently the new features. Therefore, it is necessary to design different strategies to take advantage of them.

In this work we present an operational framework to run ensemble simulations in HPC platforms. A set of experiments are presented in order to validate the suitability of this technique. Moreover, the derived impact regarding the I/O and data management aspects is analyzed.

EC-Earth, the European Community Earth System Model

EC-Earth is a seamless Earth System Model (ESM) selected by IC3 to perform climate research as end users and also as part of the EC-Earth consortium. In this study EC-Earth v3.0.1 was used and here is below a brief description of EC-Earth components:

- IFS is ECMWF's Integrated Forecast System (atmospheric component for ESM).
- NEMO the Nucleus for European Modelling of the Ocean (oceanic component for ESM) not only includes the ocean circulation but also sea-ice and biogeochemistry.
- OASIS3 a coupling software used to couple IFS and NEMO with resolutions ranging from approximately 128 to 25 kilometers.

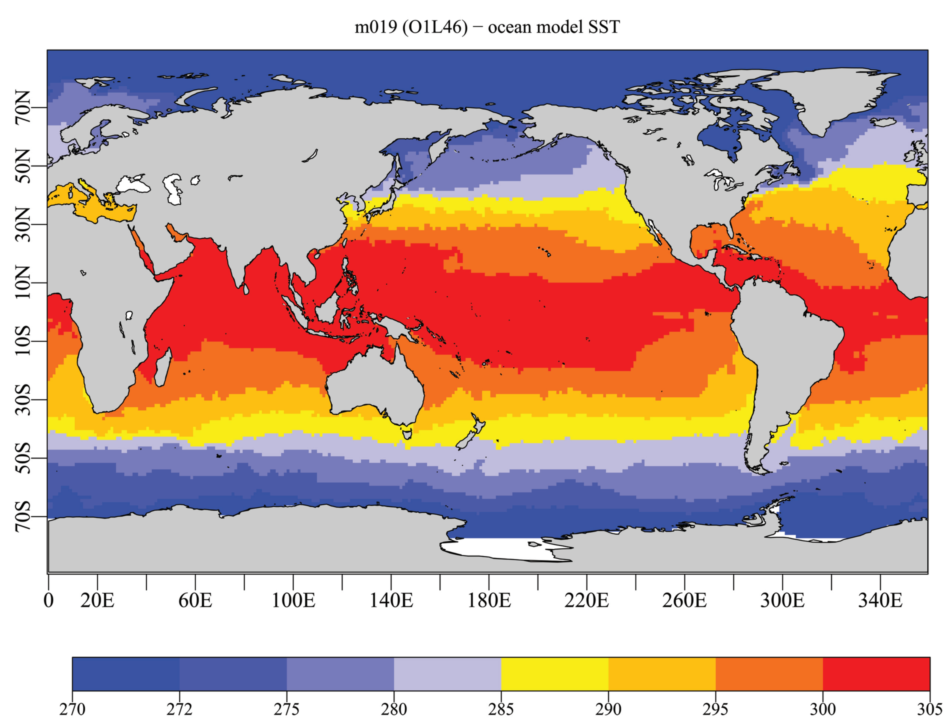

Figure 1: Example of Sea Surface Temperature forecast for one month (mean values, Kelvin scale) using the standard resolution (1 degree) of the ocean component NEMO.

Here are below the detailed list of parameters in conncection to different configurations used for this study where T159, T255 and T799 are IFS configurations and refer to approximately 125km, 80km and 25km grid-spacing respectively. On the other hand; ORCA1 and ORCA025 are NEMO configurations and refer to 110km and 25km:

Parameter	T159-ORCA1	T255-ORCA1	T799-ORCA025
IFS time step	1 hour	1 hour	12 minutes
IFS output frequency	6 hours	6 hours	6 hours
NEMO time step	1 hour	1 hour	20 minutes
NEMO output frequency	1d, 5d, 1m, 1y	1d, 5d, 1m, 1y	1d, 5d, 1m, 1y
OASIS3 coupling frequency	3 hours	3 hours	3 hours
Output size (1-month run)	0.9 GB approx.	1.6 GB approx.	21.4 GB approx.
Restart size	1.4 GB approx.	2.3 GB approx.	25 GB approx.

Table 1: Parameters of different EC-Earth3 configurations

Large Scale Climate Simulations in HPC Environments

Each configuration detailed in last section has been used in different HPC platforms. Subsequently, their main features are presented, where a) corresponds to the platform used with T159-ORCA1, b) the one used with T255-ORCA1, and c) the one used with T799-ORCA025:

a) **MareNostrum2 (Barcelona Supercomputing Center, Spain):** IBM BladeCenter JS21 with IBM Power PC 970MP processors (2.3GHz) and Myrinet interconnect technology.

b) **Lindgren (PDC Center for High-Performance Computing, Sweden):** Cray XE6 system, based on AMD 12-core Magny-Cours Opteron (2.1 GHz) processors and Cray Gemini interconnect technology.

c) **Jaguar (Oak Ridge National Laboratory, USA):** Cray XT5 system, based on AMD 6-core Istanbul Opteron (2.6 GHz) processors and InfiniBand network.

EC-Earth3 at Lindgren, PDC

Number of Start Dates		1	5	10	10	20
Number of Members		1	5	5	10	10
Number of Independent Simulations		1	25	50	100	200
T159-ORCA1	Cores	144	3600	7200	14400	28800
	Wall-clock Time (Hours) / year	5	5	5	5	5
	CPU Time (Hours) / year	720	18000	36000	72000	144000
	Output Size (GB) / year	10,80	480	960	1920	3840
T255-ORCA1	Cores	360	9000	18000	36000	72000
	Wall-clock Time (Hours) / year	5	5	5	5	5
	CPU Time (Hours) / year	1800	45000	90000	180000	360000
	Output Size (GB) / year	19,20	5184	10368	20736	41472
T799-ORCA025	Cores	1104	27600	55200	110400	220800
	Wall-clock Time (Hours) / year	40	40	40	40	40
	CPU Time (Hours) / year	44160	1104000	2208000	4416000	8832000
	Output Size (GB) / year	256,80	6420	12840	25680	51360

Figure 2: EC-Earth simulation features for each configuration at Lindgren, PDC.

A well-known goal among the climate research community is to be able to simulate 10 years per wall clock day, which means to simulate approximately 1 hour per second. In Figure 2 the detailed scalability analysis made with respect to Lindgren shows that the current community ambitions are not being met clearly and obviously there is high need to exploit HPC environment further.

Operational Framework to Run Ensemble Simulations

Due to scalability limit of the ESM; as that could also be seen in Figure 3 it seems testing and validating the option of running ensemble simulations could be helpful in explotiting the utilization of HPC power in an efficient way.

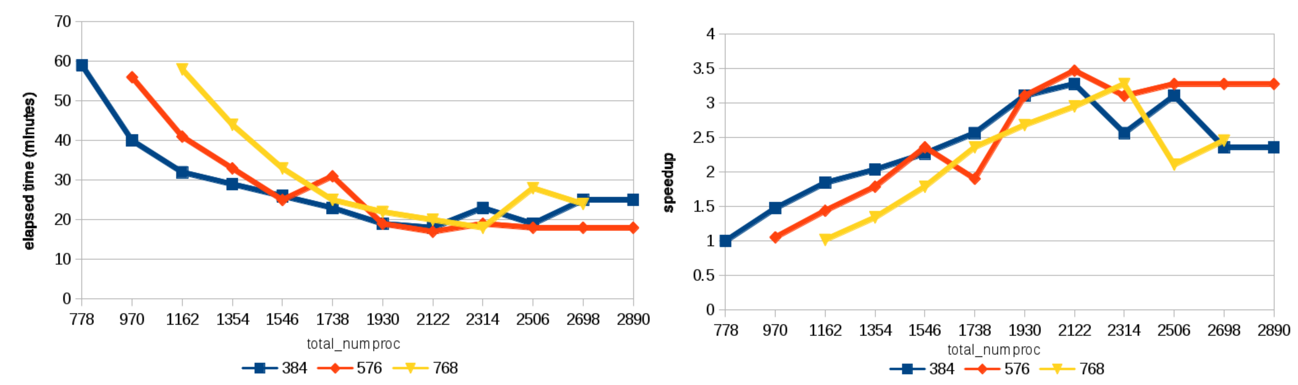

Figure 3: Left: Elapsed times of the T799-ORCA025 configuration as a function of the total number of processors for 5-day simulations, running in Jaguar. Right: Scaling of the T799-ORCA025 configuration as a function of the total number of processors for 5-day simulations. Blue, orange and yellow curves correspond to results for tests with NEMO using 384, 576 and 768 CPUs. Speedup values are relative to the reference test (778 processors).

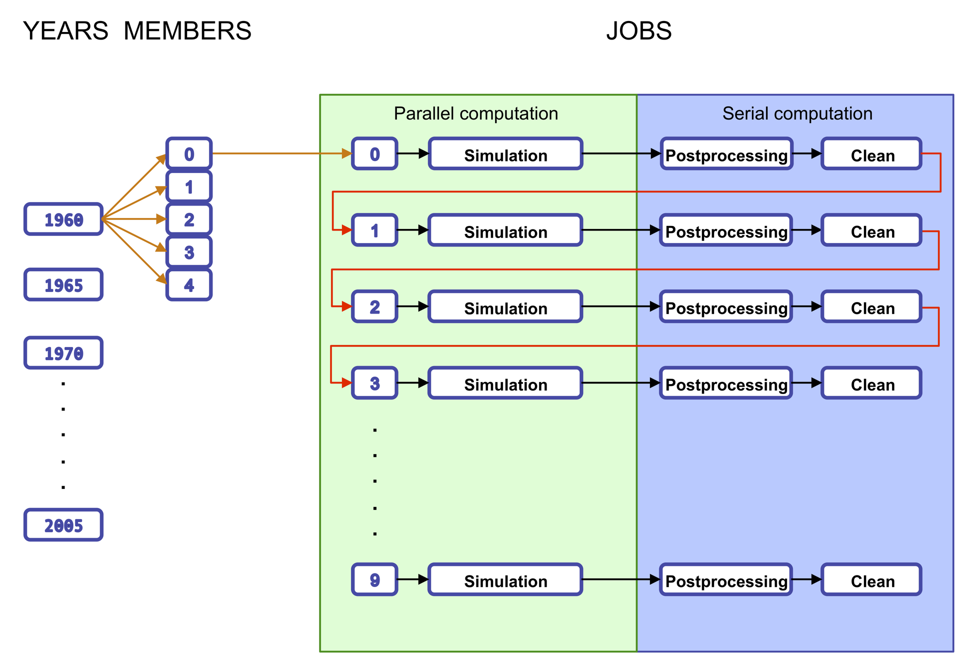

Figure 4: Sample experiment setup.

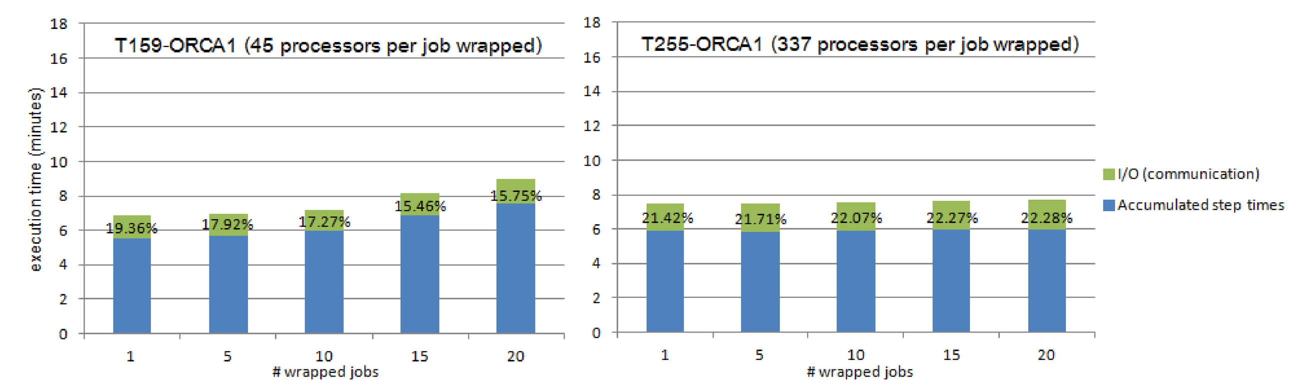

Figure 5: Average simulation times as a function of the number of jobs wrapped using configurations T159-ORCA1 (run in MareNostrum2) and T255-ORCA1 (run in Lindgren). Green portions show communication times.

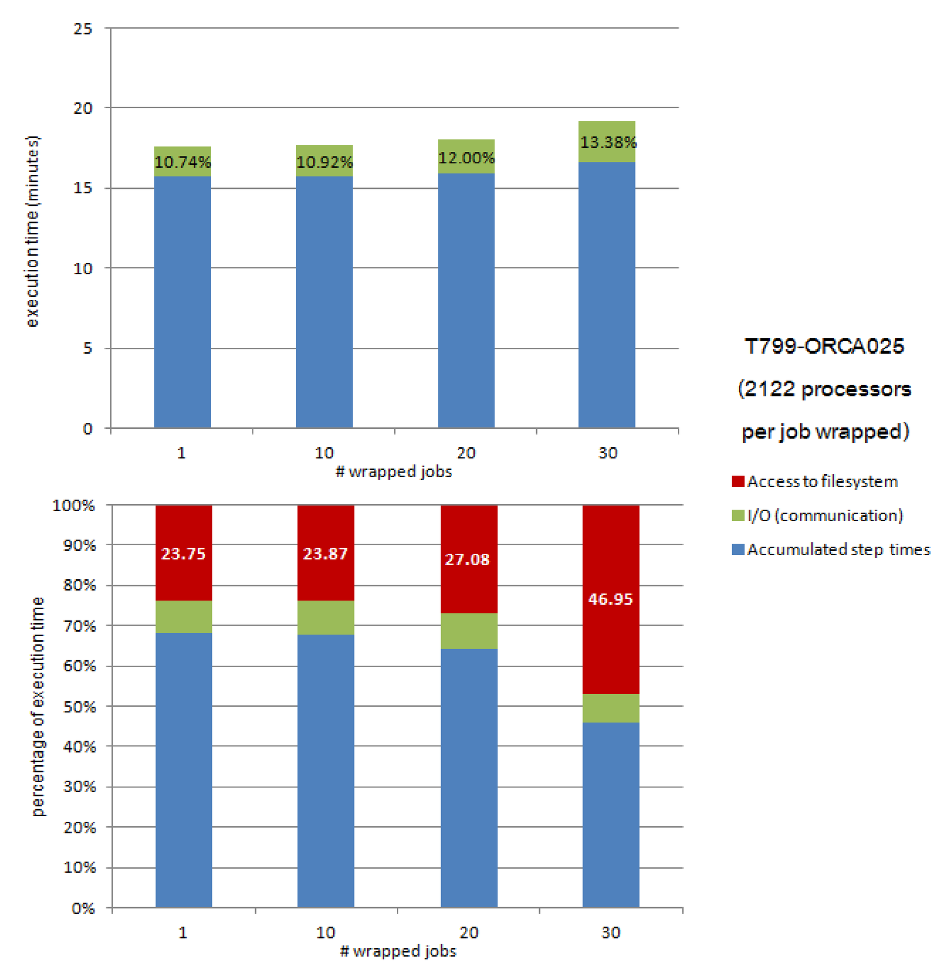

Figure 6: Top: average simulation times as a function of the number of jobs wrapped using configurations T799-ORCA025 (run in Jaguar). Green portions show communication times. Bottom: percentage of execution times, including filesystem access time

Conclusions

- EC-Earth3 presents a good scalability upto a few thousand processors only.
- A wrapping technique is exploited and validated where many independent simulations are run as a big single job.
- I/O does not exihibit the bottleneck with wrapping technique but in future it may pose issues as if the model's scalability is improved or the frequency of I/O time steps increases.
- While management of the data produced with this technique seems to be challenging.

Acknowledgements

This study was supported by the European Unions FP7-funded SPECS (308378) and IS-ENES2 (312979) projects, the MICINN-Spain under contract TIN2011-28689-C02-01 and the Catalan Government. It used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory through the grant CLI055, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.